

EUROPEAN PATENT APPLICATION

Application number: 87112735.3

Int. Cl.4: **G06F 9/46**

Date of filing: 01.09.87

Priority: 03.09.86 US 903495

Date of publication of application:
09.03.88 Bulletin 88/10

Designated Contracting States:
BE DE FR GB

Applicant: **WANG LABORATORIES INC.**
One Industrial Avenue
Lowell, MA 01851(US)

Inventor: **Cross, Charles B.**
3 Bonnie Lane
Billerica, MA 01821(US)
Inventor: **Moy, Diana Y.**
6 Old Farm Road
Wayland, MA 01778(US)

Representative: **Behrens, Dieter, Dr.-Ing. et al**
Patentanwälte **WUESTHOFF-V.**
PECHMANN-BEHRENS-GOETZ
Schweigerstrasse 2
D-8000 München 90(DE)

Multitask subscription data retrieval system.

EP 0 258 867 A2
 A multitask multiuser system provides for efficient transfer of data from a remote data base to individual subscribers and has particular utility in the distribution of data, especially of stock market data. A primary provider distributes the incoming data directly to user tasks or to an inquiry provider or a monitor provider. The inquiry provider responds to specific inquiries by users for information in the data base. The monitor provider maintains lists of information which are being monitored by the host computer for individual users. The inquiry provider and the monitor provider do not repeat requests to the remote data base where a similar request is already pending from another user. Data transfer paths between tasks are established by a code module which may be linked to any of the tasks. The transfer paths are established using information from a configuration list and they are monitored by the operating system through a wait list established for each user task. Providers in the system may establish subscriber lists through the code module.

MULTITASK SUBSCRIPTION DATA RETRIEVAL SYSTEM

Background of the Invention

Local computers are able to gain access to large volumes of information by communicating over telephone lines with remote data bases. The remote data base has storage capabilities which far exceed that which would be feasible at most local computers and can serve as a central storehouse of information.

A data base that was developed by the Walsh, Greenwood Information Systems, Inc. and which is now maintained by Wang Financial Information Services Corporation is dedicated to information relating to the stock market and other financial institutions. It contains real time trade and quote information including over-the-counter, option, commodity, futures, and fixed income data, as well as news and institutional holdings. The data base allows a subscribing computer to have access to three general classes of service: broadcast, inquiry and monitoring.

The broadcast class is that in which information is simply broadcast continuously to the user. An example is the New York Stock Exchange ticker service in which all transactions which occur on the New York Stock Exchange are transmitted to all subscribers as those transactions occur. Other broadcast services include a news headlines service which scrolls through headlines received from the Dow Jones News Service and the Reuters News Service. Also, the full news items from the Dow Jones News Service and the Reuters New Service are transmitted to subscribers to allow for scrolling through the news items as they are released.

Subscribers to the data base are also able to make specific inquiries. For example, a subscriber may send a request for a quotation on any stock item and promptly receive the current information stored at the data base for that stock item. News items of interest may also be retrieved by making requests which include specific identifier symbols which identify the information of interest.

Finally, a subconsciousness may request that the remote data base monitor all of the information which enters the data base and transmit only that data which is of particular interest to the subscriber. Again, the subscriber transmits a request to the data base which includes identifier symbols.

The Walsh Greenwood Information Systems system was designed for communication with personal computers; hence for each communications line address there has been exactly one user. Each personal computer could subscribe to a particular

set of services and pay the appropriate fee for those services. Configuration and security were handled by the network host processor sending a message to each computer on each line to indicate which services the network would allow that computer to use. This provided adequate control to permit accurate billing and accounting.

Disclosure of the Invention

Multiple user systems may have multiple terminals connected through a central computer and a single communication line to the data base host computer. Different users in this local system may subscribe to different services. In that situation, the remote data base must maintain accurate files for each of the terminals. Further, the remote data base must transmit data to all subscribing terminals along a single communications line and rely on the local multiuser system to properly distribute the data to those terminals which have subscribed and which have made specific inquiries. It is to such a multiuser system that the present invention is directed.

To properly distribute incoming data to multiple users, the local multiuser system must maintain subscription records and records of specific requests and must multiplex the incoming data to the individual users. This significant task must be performed without introducing unacceptable delays in transfer of the information from the remote data base to the individual users. The avoidance of delays is particularly critical in the case of stock market information.

It is an object of the invention as claimed with its various embodiments and further developments to provide an advantageous solution to this problem.

To provide for the aimed at rapid distribution of information, one feature of the present invention is in the handling of different classes of information. Tracking of specific monitoring requests from multiple users can become cumbersome. In accordance with one feature of the present invention, a first provider task divides the incoming data stream from the remote data base into plural data streams according to the data type corresponding to each service. Selected data streams which are in response to nonmonitoring requests are transferred directly to user tasks subscribing to the data streams. A data stream which is in response to a specific monitoring request is, however, transferred to a second provider task. That second provider task then divides the data stream in response to

the specific monitoring requests into further data streams and transfers the further data streams to the user tasks. Thus, data which is simply broadcast to the user or which is the result of a specific inquiry need not be delayed by a task which must also handle the more time-consuming distribution of data in response to monitoring requests. The second provider task may also provide the decoding necessary for a particular class of data.

Specific inquiries such as news retrievals may be handled by yet another provider task. Although that class of data does not require the extensive processing of data based on monitoring requests, the data may require sufficient additional processing to warrant a separate task. However, it is preferred that quotation inquiries be handled by the first provider task. Users are typically concerned with obtaining very rapid responses to quotation inquiries and handling of such inquiries by the first provider task is not too burdensome.

Information received from data bases must typically be displayed in a particular format established by the remote data base or by local software provided to all subscribers. Access to the received data by programs developed by the user has typically been limited and, where available, access has required particular programming expertise and effort. Such access has not been available on a real time basis. A further feature of the present invention is that data paths between the provider task and the user task are established by means of a common code module which is linked to each user task. Each data path is specific to a type of data transferred to a single user. The common code module simplifies access to the different data types to which the user has subscribed and also insulates the user program from changes in the operating system of the multitask system. Any programming changes required by changes in the operating system can be handled for all programs by modifying the common code module. By establishing a distinct data path for each type of data transferred to each user, local subscriptions can be readily entered and terminated for specific services and the user program can at any point limit its access to a particular service. Thus, the internal message traffic can be minimized.

The time during which a user program must be active while waiting for data from the data base can be minimized by establishing a wait list which is monitored by the operating system. The wait list includes a mailbox address for each type of data for each user task. Through that mailbox, the operating system notifies the user task of data being transferred to the task. The operating system speeds the distribution of data because it is not necessary for each program to independently cycle through wait routines. Rather, the operating system

waits for multiple events simultaneously and then deals with whichever one occurs first. Such capabilities are not generally available to programmers writing in higher level languages and are made available in the present system by the common code module.

In establishing the data paths between providers and subscribers, a configuration list may be established for each user task from a configuration file. Each configuration list includes the name of a provider task mailbox to which requests for each type of data are to be transferred and an indication of whether the task has access to a particular type of data. Each configuration list may also indicate whether the task is a provider of data to other tasks. Provider specific routines of the common code module are limited to provider tasks.

The common code module may also be used to establish wait lists for provider tasks. A provider task need not repeatedly check whether requests are being made from user tasks, but the operating system may simply monitor the provider task mailbox for each service. A configuration list would also be generated for each provider task so that the provider task would be able to identify each mailbox relative to each service. To facilitate distribution of data to subscribers, each provider may establish a subconsciousness list relative to each service to identify those users which have established data paths with respect to each service.

Preferably, during system start up, the first provider task is responsible for providing all other tasks in the system with the information included in the configuration list. Only the first provider task has access to configuration files which include the customer access information received from the remote data base.

A provider task such as the above mentioned second provider task for providing monitoring services should avoid making identical requests for data from the remote data base due to independent requests from different users. To that end, the provider task may receive specific data requests from user tasks and compile information correlating user tasks with specific data requests. For each data request, the provider determines whether a like request is pending for another task. Only if a like request is not pending would a specific data request be transferred to the remote data base. When the data is subsequently received from the remote data base, the provider determines from the compiled information all user tasks which have requested the received data, and the data is transferred to those user tasks. Preferably, the information is compiled by generating a user tree and a symbol tree. Each user node of the user tree points

to a list of symbols identifying the data requested by that user. Each node of the symbol tree points to the users which have requested the data represented by the symbol.

Brief Description of the Drawings

The foregoing and other objects, features and advantages of the invention will be apparent from the following more particular description of a preferred embodiment of the invention, as illustrated in the accompanying drawings.

Figure 1 is a block diagram of the software architecture of a system embodying the present invention.

Figure 2 is an illustration of the configuration list, wait list and subscriber lists stored in a provider task memory by means of an application program interface (API) common code module in the system of Figure 1.

Figure 3 is an illustration of the configuration list and wait list stored in a user task memory by means of the API.

Figure 4 illustrates a configuration file retrieved by the primary provider task through API.

Figures 5A and 5B illustrate the message format of messages transferred by means of API.

Figures 6A and 6B illustrate the data structure of a symbol tree in the monitor provider of Fig. 1.

Figures 7A and 7B illustrate the data structure of a user tree in the monitor provider.

Figure 8 illustrates an add pending queue in the monitor provider.

Description of a Preferred Embodiment

The present invention relates to a multitask, multiuser data processing system and may, for example, be implemented on a VS system from Wang Laboratories, Inc. In Figure 1, the block 20 represents the overall operating system of the multitask system. Overlayed on the operating system are a number of distinct programs which share the operating system to independently complete their respective tasks. For example, three user programs 22, 24 and 26 are illustrated. Each user program is associated with a respective computer terminal 28, 30 or 32. Within each user program, the multitask, multiuser character of the system is invisible to the user; however, for intertask communications, specific procedures must be followed to establish data paths through the operating system.

The present invention relates to the transfer of requests to and the return of data from a remote data base host computer 34. The host computer gathers real time trade and quotation information over many data lines which link the host to all exchanges and several supplemental services such as Dow Jones and Spectrum Institutional Holdings. This information includes real time stock exchange, over the counter, option, commodity and fixed income data, as well as news and institutional holdings.

All communications between the remote data base 34 and the user programs 22, 24 and 26 are processed through a primary provider task 36. This task provides the data to the subscribing users. The primary provider 36 includes conventional communications software for transferring and receiving data along telephone lines. The primary provider 36 must also make an initial determination as to what type of data is being received. Specifically, the primary provider determines the particular data service of which the data is a part.

The primary provider makes an initial distribution of the data. Most data of the broadcast class, all of which is transmitted to all subscribers of the service, is transferred directly from the primary provider to the individual subscribing users. This class of services includes the Dow Jones New Service (dj), the New York Stock Exchange Ticker (tk), News Headlines (nh) and Reuters New Service (rt). An example of the transfer paths for data from these broadcast services to user programs 24 and 26 is illustrated in Figure 1. These transfer paths are illustrated as being in only one direction because, once the data path has been established, no further requests from the user program to the remote data base are required. Other data services require more complex logic on the part of the provider. To prevent the primary provider from becoming over-burdened with distribution of data from those services, additional tasks 38 and 40 are provided.

In the present example, much of the data from the inquiry class of services is transferred by the primary provider to an inquiry provider 38. For example, with the Dow Jones New Retrieval Service (nr) the remote data base responds to specific inquiries for news items relating to particular identifiers transmitted to the data base by a user. The inquiry provider 38 initially receives those requests from the user programs, determines whether a similar request is already pending, and if no similar request is pending, transfers the request through the primary provider to the remote data base 34. All data transferred back from that service is directed by the primary provider 36 to the inquiry provider 38. The inquiry provider then examines its own records to identify the user program which

made the request and transfers the responsive data to just that user program. In the example illustrated in Figure 1, only user program 22 has subscribed to the Dow Jones News Retrieval Service, but the utility of the inquiry provider increases with increased subscriptions from other user programs and with an increased number of services of the inquiry class.

A final class of service is that in which the host computer monitors data based on a specific request for information identified by identifier symbols. The remote data base monitors all data as it is received from its sources and selects that data for which a request is pending. The data is then transferred to the requesting user. Thus, the monitor class of services keeps an inquiry open. The monitor provider receives requests from the user programs and, as did the inquiry provider, only forwards requests to the remote data base which are not already pending. The remote data base, for the most part, sees the multiuser system as a single user and only transmits a particular item of data once to the multiuser system. When that data is received, it is directed to the monitor provider. The monitor provider searches its records to determine which user program or programs requested the data and transfers the data to those user programs.

Market monitor data is encoded in a compressed format. The primary provider determines only that the data being received is of the market monitoring format and transfers all of that data (MM) in coded format to the monitor provider 40. The monitor provider then decodes the data and directs it to the requesting user programs as one or more of three services. The basic market monitoring service (mm) responds to each price change in that particular stock. A select ticker service (sl) responds to each trade of the stock and identifies the volume and price of that trade. The block trade ticker (bt) identifies all trades which exceed ten thousand shares. Block trade is actually a broadcast class of service but it is encoded by the host computer with the market monitor service, so it is most efficiently decoded by the monitor provider. Further, it is expected that block trading in the future will be a more selective service. As illustrated in Figure 1, although all three services rely on the same data stream received from the remote data base, user programs may individually subscribe to select ones of the three.

One service which is of the inquiry class but which is handled entirely by the primary provider is that of stock price quotations (qt). Users expect that price quotations will be received very promptly, and the present system facilitates the prompt return of quotation data by minimizing the number of transfers of the quotation requests and return

data through the operating system. The two transfers in each direction which result from use of the inquiry provider is avoided. The primary provider transmits the stock quote request with a sequence code which identifies the request in the primary provider. The sequence code is stored with the identity of the requesting user program in a first-in/first-out storage. The remote data base returns the data in the same order in which it is requested and returns the same sequence code with the data. The primary provider then returns the data to the requesting user task as it is identified from the first-in/first-out storage so long as the sequence numbers match. If a sequence number is skipped in the return data, the primary provider is able to either make the request again or notify the user program that the request should be made again.

A basic user program might allow the user to make specific requests within a service to which it subscribes and to suitably display the returned data. However, it is contemplated that users will want to develop their own programs which are particularly suited to their particular needs. In fact, users may wish to develop programs which have as only a small part thereof the need to obtain stock market information. A pension fund management program might be an example of such a program. In the past, retrieval of information from stock market services and other data bases has offered little flexibility to the programmer and has made linking of custom programs to the remote data base difficult or even impossible. The present system facilitates such communications by means of an application program interface (API) which facilitates the establishment of data transfer paths between tasks by making full use of the operating system capabilities while minimizing the programming effort required of the user program.

API is a code module of subroutines which may be linked to any one of the user and provider tasks. Once the API code module has been linked to a user program, the program need only make simple calls to the subroutines to establish the data transfer paths. The user need not be aware of the capabilities and requirements of the underlying operating system because those capabilities and requirements have been taken into consideration in development of the code module. Further, the user program is insulated from changes in the operating system. The system designer will make appropriate changes in the API code module with changes in the operating system. Although the API code module offers its greatest advantages in interfacing user programs to the operating system, it is also used to interface the provider background tasks.

With this approach, the system designer, in making changes to the operating system, again need only modify the API code module and need not modify the background provider tasks.

A further advantage of the particular API code module to be described below is that it makes full use of the intertask message (ITM) capabilities of the Wang VS systems in creating mailboxes and transporting messages to those mailboxes. In the Wang VS systems, a task is able to list ports, or mailboxes, in an operating system plist and then "put itself to sleep". In the present system, those plists are wait lists created by the API code module. The operating system monitors the plist and, when a message is received at a port on the plist, the operating system notifies the task associated with that plist. As a result, the task program need not continue to operate as the program waits for an event to occur. This feature of the operating system is particularly useful to the provider tasks in monitoring data requests from user programs and to user programs in monitoring the receipt of data from the host computer. Such events may occur at any time.

Intertask communications are handled efficiently by means of the API code module by the establishment of three types of lists. Figures 2 and 3 illustrate example lists for the primary provider and user2. Each task which communicates through the API code module establishes a configuration list and a wait list through API routines. In addition, each provider task may establish a subscriber list for each service that it provides.

Before establishing any data transfer paths, a provider or user task must first obtain a configuration list through the primary provider. The configuration list includes a service code for each available service. The configuration list also indicates, for each service, whether the particular task is a provider of that service (indicated by a P in the list). If the task is not a provider the list indicates whether the task has access to the service (indicated by Y for yes or an N for no). Each task is also provided with an ITM port, or mailbox address, to which messages requesting that service should be transferred.

Before a user can complete a request to a provider, the provider must have retrieved its configuration list. Based on its access and provider mailbox information provided in that list, the provider must have established itself to receive requests for the particular service by listing the mailbox at which it is to receive requests on its operating system wait list. In Fig. 2, the active status indicated for each service in the wait list indicates that the primary provider is ready to receive requests for each service on the wait list. As already noted, the wait list is an operating system plist in

the Wang VS system architecture. Once the wait list has been established through the API code module, the provider program need not continue to monitor its mailboxes for the respective services. Rather, this chore is handled by the operating system by means of the wait list.

The provider may establish a subscriber list for each service and include in its configuration list a pointer to the starting address of each subscriber list (indicated by an * in Figure 2).

The user also establishes its wait list in the operating system through API routines. The user may establish its own user port for each service to which it has access. That user port then serves as the mailbox for any data returned from a provider. The wait list includes each of its mailboxes on which it expects to receive data. Each mailbox on the wait list may be set at either an active or an ignore status; the operating system will only notify a task of transferred data to those mailboxes which are active.

To subscribe to a service, a user looks to its configuration list in an API routine to obtain a provider mailbox address. Through the routine, it then constructs a subscription message as illustrated in Figure 5B and transmits that message to the designated provider mailbox. It may then wait to be notified by the operating system through its user port that a message has been received. The subscription message is transferred by the operating system into the provider task mailbox and the provider task is notified that a message has been received. The provider then processes the subscription request and may place the user's return mailbox and other data for that service in the provider's subscriber list for that service. That other data includes the subscriber's workstation number, user I.D. and key.

By defining in the configuration list a unique provider mailbox for each service, the system allows each user program to independently establish data paths for the several services to which it subscribes. Although the host processor may include the particular user program as a customer and thus grant access to data for a particular service, within the local system any user may fail to subscribe or temporarily cancel any subscription to any service at any time without affecting subscriptions to other services. As a result, internal message traffic through the operating system can be minimized. The host processor is not informed of the user's status as a subscriber within the local system, so subscriptions and cancellations within the system do not affect the customer status of the user.

To provide for prompt access to the configuration information, a distinct configuration list is generated for each task in the system. Each task is able to efficiently obtain access to its list without requiring further transfer of data from a configuration file through the operating system. However, in a multitask system in which a global memory is readily available to all tasks, the individual configuration lists may not be required. Rather, the information may be obtained from a centralized file. It is important, however, that each task have ready access to the provider mailbox addresses designated by the system and to its own access information with respect to each available service.

A more detailed description of the creation of the configuration, wait and subscriber lists and the API subroutines which allow for the efficient transfer of data follows.

As the initial receiver of all information from the remote data base, the primary provider 36 has been selected as the focal program for initialization of programs with respect to transfers from the remote data base. As a part of the primary provider's initialization processing, it makes a call to a configuration initialization API routine called CNFINIT. CNFINIT reads a configuration file from disc storage. The primary provider also starts the communication line. When the host computer 34 detects that the line is active, it responds by sending out a services list for customers as discussed below. The primary provider consolidates the information received from the remote data base and from the disc storage to create the configuration file of Fig. 4.

The first item of the configuration file is a key for each task to be involved in the data transfer. In effect, the key is a logical name for a configuration record for each task. The keys PRIPRO, INQPRO and MONPRO are the names provided for the primary provider, inquiry provider and monitor provider, respectively. A further key ADMIN allows for the implementation of an administrative program if it is required for the particular needs of a system. Such a program may control designation of keys and control access to services by particular tasks and particular workstations.

The programs associated with the first four keys of the file are provider programs which typically operate in background. An additional three keys USER1, USER2 and USER3 are indicated for each of the three user programs shown in Fig. 1. Of course, the system is not limited to three user programs. These programs are typically subscribers and may or may not operate in foreground.

The terminal I.D. field is the identifier by which the remote data base knows each user on a system. The data base grants rights to particular services and bills for those services in a per terminal

identification. Only user tasks need terminal I.D.s since they are the only true consumers of data. Suppliers may be thought of as merely an extension of the remote data base and require no terminal I.D. The terminal I.D. also serves as an alternate key; the remote data base is unaware of the first key and sends its configuration records tagged by the terminal I.D. The terminal I.D.s are assigned by the remote data base when a customer requests new or expanded services.

The task number is forwarded to the primary provider with a key as each task is initialized. The status field indicates whether a key is in use and, if so, the work station number at which the associated task is being used. The status field may indicate that a task is being operated in background. The workstations fields may be used by an administrator program to restrict the workstations on which a given key and its corresponding privileges may be used. The key may be restricted from or to a list of workstations which is simply illustrated by an asterisk in Fig. 4.

The services field is expanded in the lower portion of Fig. 4. This list of fields defines, for each key, the associated program's role with respect to each defined service. The program may have any of three roles in connection with a given service: it may have full subscription privileges (Y); it may be the provider of the service (P); or it may have no rights at all for the service (N). The subscriber rights given by the remote data base may be overruled by a local administrator program as indicated by y, for subscription allowed, and n, for subscription denied. The roles of providers are defined by the system designer.

During CNFINIT, primary provider 36 also obtains a provider mailbox file from disc storage. The provider mailbox for each service is later included in each configuration list of each task as illustrated in Figs. 2 and 3. Using the information from the provider mailbox file and the configuration file, the primary provider builds a prototype of the configuration list to be used by each task employing the API code module. The prototype configuration list contains one record for each defined service, and each record includes the mailbox, which is the ITM port name, of the provider of the service.

With the configuration file and the prototype configuration list on hand, the primary provider then generates its own personal configuration list by calling SCLINIT. By providing its key PRIPRO, it obtains its configuration list, based on the prototype list, with its access code P, Y or N taken from the configuration file for each service. Where a P is indicated for a service in the configuration list, the provider establishes as its mailbox the provider mailbox indicated in the configuration list. The provider also establishes a subscriber list for each

such service and sets the current list size to zero. Further, each mailbox established by the provider program is placed by the API subroutine on an operating system plist for that task which will serve as a wait list. The wait status for each provider mailbox on the wait list is set to active. Finally, the primary provider program calls an API routine SCLWAIT by which it signals the operating system that it will wait for an indication that a message has been received at any of its active mailboxes on the wait list.

With the initialization of the primary provider, other tasks may be initialized. To that end, each task calls the API routine SCLINIT, whereby it establishes an operating system wait list. As illustrated in Fig. 3, a wait list may initially include three pseudoservices emergency (em), workstation (ws) and timer (tm). The emergency port allows the task to receive emergency messages at a high priority on the wait list. The workstation service allows the user program to utilize the wait list established by the API code module to wait for events which are to occur at the workstation terminal. With this pseudoservice, the user program need not delay receipt of data while waiting for an event such as a key stroke to occur at the terminal. Finally, the timer pseudoservice allows the program to set a timer by way of the wait list. It allows the user program to limit the time that it will wait for other events to occur.

During the SCLINIT routine, the task also obtains its configuration list which will provide it with the information required to subscribe to the services in the system. It obtains that configuration list from the primary provider task and, to that end, the API subroutine has a default miniconfiguration list by which every task is configured as a user of configuration data. A user subscribes to the configuration data through a mailbox "cfcf" to the primary provider.

When the configuration subscription message arrives at the mailbox "cfcf" of the primary provider, the operating system awakens the primary provider in the SCLWAIT routine. In response to the subscription message for configuration data at its mailbox "cfcf", the primary provider uses the user key provided in the subscription message to obtain the appropriate record in the configuration file and mark a copy of the prototype configuration list with the appropriate access code for each service for that key. The primary provider returns the customized configuration list to the subscribing mailbox, for example 20cf, and updates its master configuration file to indicate that the particular key is in use.

Through individual use of the SCLINIT routine, each task may have a full configuration list and a wait list which is empty but for the emergency, workstation and timer pseudoservices. The data paths illustrated in Fig. 1 have not yet been established. To establish the paths, each provider task calls a routine PRVOPEN, by which it establishes a subscriber list which is initially empty. Both provider tasks, automatically during PRVOPEN, and subscriber tasks call a routine SCLOPEN. By the routine SCLOPEN, a provider task sets its mailbox for the service being opened to that indicated in the configuration list, and a subscriber task sets the mailbox to its task number plus the service code, as for example 20qt in Fig. 3. Any future messages to be received by the particular task for that service are received at the designated mailbox.

Some tasks may serve as a subscriber to some services and as a provider of others. For example, the monitor provider serves as a subscriber to market monitoring service MM and as a provider of market monitoring service mm, select ticker service sl and block trade service bt. Similarly, the inquiry provider is a subscriber to news retrieval service NR and a provider of the distributed news retrieval service nr. Separate SCLOPEN routines must be called for each service.

In the SCLOPEN routine, the mailbox is placed on the wait list for that task but is set to "ignore". A provider task then calls SCLLISTEN, whereby the particular mailbox on the wait list is set to active status. A subscriber task will typically first call a routine SCLSUBSCRIBE. By that routine a subscriber builds a subscription message as illustrated in Fig. 5B and transmits that subscription message to the provider mailbox obtained from its configuration list. The subscription message includes a return mailbox for the service at the subscriber and a flag indicating whether the message is for a subscription or cancellation. The message also includes a task number, a user I.D. and the user's key. The provider responds to the message by adding the subscriber to the subscriber list for that service. The subscriber then typically enters SCLLISTEN whereby the associated mailbox on the wait list is set to active status.

Data returned to a subscriber mailbox by a provider may be read by entering the API routine SCLREAD in which the user obtains whatever information is then in the mailbox. Usually, however, a user enters the routine SCLWAIT whereby it signals the operating system to notify it when an event occurs at any of its active mailboxes.

Once a user has subscribed to a service, it may make specific requests of the remote data base through its provider by an SCLSEND routine in the API code module. Using this routine, the user program generates a message as illustrated in

Fig. 5A which includes its mailbox for the service and a flag and message appropriate for the particular service and request. Under SCLSEND, the provider mailbox is obtained from the subscriber's configuration list and the message is sent. The message is processed by the provider and, in the case of an inquiry provider or monitor provider, the provider might then generate its own message to be transmitted to the primary provider and on to a remote data base. The messages to the primary provider would be sent over paths previously established through the SCLOPEN and SCLSUBSCRIBE routines.

Finally, when data is received by the provider from the host processor 34, the provider must forward the data to its subscribing users. It may call a PRVSEND routine whereby the provider refers to its subscription list for the particular service and sends the incoming data to those subscribers. Alternatively, a provider such as the monitor provider may determine the mailboxes to which the data must be transmitted through its own internal lists.

In addition to the routines thus far described, there are inverse routines such as SCLTERM which is the inverse of SCLINIT, SCLCANCEL which is the inverse of SCLSUBSCRIBE, SCLIGNORE which is the inverse of SCLLISTEN and SCLCLOSE which is the inverse of SCLOPEN.

An additional function SCLLIST allows a user program to add a number of user specified events to the wait list used by SCLWAIT. This pseudoservice (us)-like the workstation pseudoservice, allows the user program to avoid further wait processing during which the receipt of incoming data might be delayed.

A more detailed description of the monitor provider task will now be presented with reference to Figures 6 through 8. The monitor provider must not only maintain a list of subscribers for each of the market monitor, select ticker and block trading services; it must also transfer specific requests that particular market symbols be monitored by the host and correlate the subscriber list with the requested symbols in order that data received from the host can be properly distributed. To that end, the monitor provider relies on a symbol tree illustrated in Fig. 6A and a user tree illustrated in Fig. 7A. Each tree is a binary threaded tree which additionally includes pointers to linked lists.

In the case of the symbol tree, each node of the tree identifies a symbol which serves as an identifier to a particular security which is being monitored. Each node of the tree points to a linked list of users which have requested information with respect to that symbol. In the user tree, each node of the tree identifies a subscriber to any of the services from the monitor provider, and each node

may point to a linked list of symbols which are being monitored for that user. For each symbol in the list linked to each user node, there is a corresponding symbol node in the symbol tree.

The user tree grows as additional user tasks subscribe to the monitor provider services. The symbol tree grows as the number of specific requests within each service increases. In each case, the initial user to subscribe or the initial symbol to be requested serves as the root of the tree. Thereafter, additional nodes are added to the left or right of a preceding node depending on whether the new node has a lower or higher alphabetic value. Thus, any symbol which alphabetically precedes the symbol MMM.N at the root of the symbol tree will be found to the left of the tree. With the use of a binary tree, the time required to search the list for a particular item is reduced substantially. This search time is particularly important where as many as 20 users and as many as 50 symbols per user may be listed in the trees.

Each node of the symbol tree includes the information listed in Fig. 6B. Each node includes the security symbol, the exchange on which that security is traded, and the most recent price for that security received from the host. It further includes pointers to the end users on the user list linked to that node. Pointers to and within each linked list are illustrated by broken arrows in Fig. 6A. The linked lists are doubled linked to facilitate adding and deleting nodes from the list. Pointers to the left and right children of each node are provided along with a pointer to the parent of the node. These pointers are illustrated by solid arrows in Fig. 6A. Threading of the tree by the pointers to the parents simplifies restructuring of the tree when symbols are added or removed. Each user node of the linked list includes the service by which the particular user has requested the symbol. A service minding table included in each symbol node identifies all of the services included within its linked user list.

Each user node of the user list carries the information presented in Fig. 7B. It includes the return mailboxes identified in the subscriber's subscription message under the routine SCLSUBSCRIBE, the user's terminal I.D. and the user's key. Further, each node includes pointers to the symbol list associated with that node and pointers to the children and parent of the node. Further, the user node includes its own service minding table. This minding table lists all of the services to which the particular user has subscribed. As illustrated by user 7 in Fig. 7A, the particular user need not have specific symbol requests pending to have a user node.

When a user subscribes to a service, through SCLSUBSCRIBE, the monitor provider scans the user tree to determine whether a user node has already been established for that task. If so, the new service to which the task is subscribing is added to the service minding table. If not, a new node is added as a leaf to the tree.

Thereafter, by means of an SCLSEND routine, the task may request that a particular symbol be monitored for that service. Prior to adding the symbol to the tree, a request must be transmitted to the host processor and the acknowledgement in the form of a price must be received from the host. A request is only transmitted to the host processor if the same symbol within the same service has not previously been requested. To make that determination, the monitor provider first scans the symbol tree for the symbol which is to be requested. If it locates that symbol, it notes the service minding table to determine whether the symbol has been requested for the desired service. If the symbol and service are located in the symbol tree, the user is added to the linked user list to the symbol node, and the symbol is added to the linked symbol list of the user node. If, however, the symbol and service are not located on the symbol tree a request is made to the host.

When the request is made to the host, it is given a sequence number which is transmitted with the request. The host will acknowledge the request on a first-in/first-out basis and return the sequence number. To determine whether an acknowledgement is received, a requested symbol is placed in an add pending queue illustrated in Fig. 8. Each node of the queue points to the next succeeding symbol to be requested. The sequence number is included with the requested symbol. In addition, each node of the queue points to a user node which may be included in a list.

When an acknowledgement is received, the symbol node is removed from the queue, and the pointer to the user node is transferred to that symbol's node in the symbol tree. If the symbol does not already have a node in the tree, a node is established. Further, a node in a linked symbol list in the user tree is added for the particular user.

If an acknowledgement sequence number is received out of order, the system assumes that the skipped request was not successful and transmits a new request for the same symbol. The symbol is placed at the end of the add pending queue with a new sequence number. Also, a counter is clocked. The counter allows a request to be retransmitted only three times to avoid an endless loop of requests for a symbol which is not being acknowledged.

When a user requests that a symbol be deleted from its service, the user tree is scanned to locate the user and that symbol is removed from the linked list. Then, the symbol tree is scanned for the symbol and the user is removed from its linked list. Further, as the linked list is scanned to locate the user, it is determined whether that user is the last to include that symbol for a particular service. If so, that service is deleted from the symbol node minding table. Thereafter, in scanning the symbol tree to determine whether a particular service is pending for a symbol, it is not necessary to scan the linked user list because the minding table has been updated. If the deleted user is the last node on the linked user list, the full node is removed from the symbol tree. Also, a delete request is forwarded to the host processor.

If a delete request is not properly received and processed by the host, the system may in the future receive update information which the host expects to be distributed. On scanning the symbol tree to locate the users of that data, no corresponding symbol node will be located. In that event, a delete request is again transmitted to the host processor.

If the symbol is located in the symbol tree when information is received from the host, the linked user list is then scanned to identify the users which are to receive the particular data. The data is forwarded to those users at their appropriate mailboxes determined by scanning the user tree.

The user tree is required in order to clear the symbol tree of symbols associated with that user when the user program calls the SCLCANCEL routine, forwarding a cancellation message to the monitor provider. On receipt of the message from that routine, the user tree is scanned for the user and the symbols are identified by scanning the linked symbol list. Then, the symbol tree is scanned for those symbols and the user is removed from the linked user list. Again, if the user is the last user to be linked to that symbol, the symbol is removed from the symbol tree and a delete request is forwarded to the host.

It can be noted that the block trade symbol bt is not included in either of the sample trees. This is because the block trade service is a broadcast service. It is only handled by the monitor provider because it is encoded with the market monitor data. However, as a broadcast service, block trading is readily handled by the API routine PRVSUBSCRIBE by which the user is added to a subscriber list and PRVSEND by which data is distributed to the listed users.

While this invention has been particularly shown and described with references to a preferred embodiment thereof, it will be understood by those skilled in the art that various changes in form and

details may be made therein without departing from the spirit and scope of the invention as defined by the appended claims. For example, although the invention is directed primarily to a multitask system in which the tasks are performed in a single processing unit, the tasks may be distributed to other processing units while still incorporating certain features of the invention. Further, the invention has application to nonmarket data.

Claims

1. In a multitask data processing system having means for receiving data of different types in response to nonmonitoring requests for data and specific monitoring requests for data, the specific monitoring requests requiring the remote data base to monitor data base sources with respect to specific identifiers, a method for distributing the data to plural user tasks subscribing to the data comprising:

in a first provider task dividing a data stream of plural data types from the remote data base into plural data streams according to the data types; transferring from the first provider task at least one selected data stream in response to nonmonitoring requests directly to user tasks subscribing to the data streams and transferring at least one selected data stream in response to specific monitoring requests to at least a second provider task; and in the second provider task, dividing the data stream in response to specific monitoring requests into further data streams and transferring the further data streams to user tasks requesting the data streams.

2. In a multitask data processing system having means for receiving data of different types, a method of distributing the data from a provider task to plural user tasks subscribing to the data, the method comprising:

by means of a common code module linked to each user task, establishing data paths between the provider task and the user tasks, each data path being specific to a type of data transferred to a single user; and

by means of the provider task, dividing a data stream from the remote data base into plural data streams according to the data type and transferring the plural data streams to the user tasks through the established data paths.

3. In a multitask data processing system having means for receiving data of different types a method of distributing the data from a provider task to plural user tasks subscribing to the data, the method comprising:

in response to calls from the user tasks, establishing data paths between the provider task and the

user tasks, each data path being specific to a type of data transferred to a single user; and by means of the provider task, dividing a data stream from the remote data base into plural data streams according to the data type and transferring the plural data streams to the user tasks through the established data paths.

4. In a multitask data processing system having means for receiving data, a method of distributing the data to plural user tasks comprising in a provider tasks:

receiving specific data requests from the user tasks and compiling information correlating user tasks with specific data requests;

for each data request, determining whether a like request is pending for another task;

only if a like request is not pending, transferring to the remote data base the specific data request;

receiving specific data from the remote data base;

determining from the compiled information all user tasks which have requested the received specific data; and

transferring the specific data to user tasks which have requested the data.

5. In a multitask data processing system having means for receiving data of different types, a method of distributing the data from provider task to plural user tasks subscribing to the data, the method comprising:

by means of a common code module linked to each user task, establishing data paths between the provider tasks and the user tasks, each data path being specific to a type of data being transferred to a single user, the data paths being established by each user task using a configuration list for each user task from a configuration file, each configuration list including the name of a provider task mailbox to which requests for data of each type are to be transferred and an indication of whether the task has access to a particular type of data;

monitoring, by means of a wait list established in an operating system mailboxes with respect to each type of data to the provider task to identify request from user tasks and to user tasks to identify data being transferred to the user tasks;

receiving in a monitor provider task specific data requests from the user tasks and compiling information correlating user tasks with specific data requests, determining whether a like request is pending for another task, and only if a like request is not pending transferring through a primary provider task to the remote data base the specific data request;

in the primary provider task, dividing a data stream of plural data types from the remote data base into plural data streams according to the data types; transferring to nonmonitoring requests directly to

user tasks subscribing to the data streams and transferring at least one selected data stream in response to specific monitoring requests to the monitor provider task; and

in the monitor provider task, dividing the data stream in response to specific monitoring requests into further data streams and transferring the further data streams to user tasks requesting the data streams.

6. A method as claimed in claim 1 further comprising the step of decoding the data stream in the second provider task.

7. A method as claimed in claim 1 further comprising, by means of a common code module linked to each user task, establishing data paths between the provider tasks and the user tasks, each data transfer path being specific to a type of data transferred to a single user task.

8. A method as claimed in claim 7 further comprising the step of generating in at least the first provider task, a subscribers list for identifying, relative to each type of data, the tasks which are to receive the type of data.

9. A method as claimed in claim 1 further comprising the step of monitoring.

10. A method as claimed in claim 1 wherein the second provider task receives specific data requests from the user tasks and compiles information correlating user tasks with specific data requests, the data being compiled by generating a user tree and a data-type tree, each user node of the user tree pointing to a list of data types requested by that user and each node of the data-type tree pointing to the users which have requested that type of data.

11. A method as claimed in claim 1, 3, 7 or 9 further comprising, in establishing data paths, the step of generating a configuration list for each of the provider tasks and the user tasks from a configuration file, each configuration list including the name of a provider task mailbox to which requests for data of each type are to be transferred and an indication of whether the task has access to a particular type of data.

12. A method as claimed in claim 1, 3, 7 or 11 further comprising the step of monitoring, by means of a wait list established in an operating system, mailboxes with respect to each type of data to the provider tasks to identify requests from user tasks and to user tasks to identify data being transferred to the user tasks.

13. A method as claimed in any of claims 1 to 5 wherein the data comprises stock market information.

14. A method as claimed in claim 13 further comprising transferring stock quotation requests from the user tasks to the first provider task, providing each such request with a sequence number

in the first provider task and transferring the request to the remote data base, the first provider task maintaining a record of users requesting the data with respect to sequence numbers and, by means of sequence numbers returned with the data, confirming receipt of data responses to the requests and forwarding the received data to the users which made the quotation requests.

15. A method as claimed in claim 2 further comprising the step of monitoring, by means of a wait list established through the common code module in an operating system, a mailbox to a user task with respect to each type of data to notify the user task of data being transferred to the user task.

16. A method as claimed in claim 2 wherein the data paths are established by routines of the common code module based on system configuration information available to the common code module routine, the routine being called by a user task without regard for the provider task architecture.

17. A method as claimed in claim 16 wherein the configuration information is available through a configuration list which includes the name of a provider task mailbox to which requests for data of each type are to be transferred and an indication of whether the task has access to a particular type of data.

18. A method as claimed in claim 2 wherein data streams of a first type are transferred to specific user tasks depending on the content of the data streams, and data streams of a second type are transferred to all user tasks from a list of user tasks.

19. A method as claimed in claim 15 further comprising the step of including in the wait list events specified by the user other than data transfers from a provider task.

20. A method as claimed in claim 2 or 15 further comprising, in establishing data paths, the step of generating a configuration list for each user task from a configuration file, each configuration list including the name of a provider task mailbox to which requests for data of each type are to be transferred and an indication of whether the task has access to a particular type of data.

21. A method as claimed in claim 20 wherein each configuration list indicates whether the task is a provider of data to other tasks, the method further comprising the step of limiting provider specific routines of the common code module to provider tasks.

22. A method as claimed in claim 15 or 20 further comprising the step of monitoring, by means of a wait list established through the common code module in an operating system, a mail-

box to a provider task with respect to each type of data to notify the provider task of data being transferred to the provider task.

23. A method as claimed in claim 22 further comprising the step of generating a configuration list for each provider task from a configuration file, each configuration list including the name of a provider task mailbox to which requests for data of each type are to be transferred. 5

24. A method as claimed in claim 2 or 22 further comprising the step of generating in at least the provider task, a subscribers list for identifying, relative to each type of data, the tasks which are to receive that type of data. 10

25. A method as claimed in claim 4 wherein the information is compiled by generating a user tree and a symbol tree, each user node of the user tree pointing to a list of symbols identifying data requested by that user and each node of the symbol tree pointing to the users which have requested the data represented by the symbol. 15 20

26. A method as claimed in claim 25 wherein each node of the user tree includes a listing of services to which that user has subscribed.

27. A method as claimed in claim 25 wherein each node of the symbol tree includes a listing of services through which users have made requests for the data identified by the symbol. 25

28. In a data processing system, a data structure comprising: a first tree structure (Fig. 6A) of data of a first data type, each node of the first tree structure able to point to a list referencing data of a second data type; and 30

a second data type, each node of the second tree structure able to point to a list referencing data of the first data type. 35


29. A data structure as claimed in claim 28 wherein the data of the first data type identifies services, and the data of the second data type identifies users of the services. 40

30. A data structure as claimed in claim 29 wherein a list referencing data of the second data type is a list referencing all users of a service, and a list referencing data of the second data type is a list referencing all services used by a user. 45

50

55

Fig. 1



PRIMARY PROVIDER

CONFIGURATION LIST			WAIT LIST	
SERVICE	ACCESS	PROV. MAILBOX	PROV. MAILBOX	WAIT STATUS
cf	P	cfcf	l2em	A
MM	P	MMMM	cfcf	A
NR	P	NRNR	MMMM	A
ad	Y	adad	NRNR	A
qt	P	qt qt	adad	A
Otk	P	tk tk	qt qt	A
dj	P	dj dj	tk tk	A
rt	P	rt rt	dj dj	A
nh	P	nhnh	rt rt	A
bt	N	bt bt	nhnh	A
sl	N	sl sl	l2tm	I
mm	N	mmmm		
nr	N	nrnr		
SUBSCRIBER LISTS				
cf	MM	NR	qt	nh
13cf	13MM	14NR	20qt	20nh
14cf			20djt	20rt
20cf			22tk	22rt
, 22cf				

Fig. 2

USER.2

CONFIGURATION LIST					WAIT LIST	
<u>SERVICE</u>	<u>ACCESS</u>	<u>PROV. MAILBOX</u>	<u>USER MAILBOX</u>		<u>USER MAILBOX</u>	<u>WAIT STATUS</u>
cf	Y	cfcf	20cf		20em	A
MM	N	MMMM			20qt	I
NR	N	NRNR			20tk	A
ad	N	adad			20dj	I
qt	Y	qtqt	20qt		20rt	I
Otk	Y	tktk	20tk		20nh	A
dj	Y	djdj	20dj		20mm	A
rt	Y	rtrt	20rt		<u>OTHER EVENTS</u>	
nh	Y	nhnh	20nh		ws	A
bt	N	btbt			us	A
sl	N	slsl			tm	I
mm	Y	mmmm	20mm			
nr	N	nrnr				

Fig. 3

Key	Terminal ID	Task No.	Status	Workstations restricted	list	Services list
PRIPRO	00000	012	BG	from	*	*
INQPRO	00000	014	BG	from	*	*
MONPRO	00000	013	BG	from	*	*
ADMIN	00000	029	014	to	*	*
USER1	A1074	019	023	to	*	*
USER2	A1075	020	024	from	*	*
USER3	A1077	022	BG	to	*	*

Key	Services list																
	internal				services available to users												
	cf	MM	NR	ad	qt	tk	dj	rt	nh	bt	sl	mm	nr				
PRIPRO	P	P	P	Y	P	P	P	P	P	N	N	N	N				
INQPRO	Y	N	Y	Y	N	N	N	N	N	N	N	N	P				
MONPRO	Y	Y	N	Y	N	N	N	N	N	P	P	P	N				
ADMIN	Y	Y	Y	P	Y	Y	Y	Y	Y	Y	Y	Y	Y				
USER1	Y	N	N	N	Yn	Yn	N	N	N	Yy	Yy	Yy	Yy				
USER2	Y	N	N	N	Yy	Yy	Yy	Yy	Yy	N	N	Yy	N				
USER3	Y	N	N	N	N	Yy	Yy	Yy	N	Yn	N	Yn	Yn				

Fig. 4

Description	Type	Length	Remarks
return mailbox	text	4	sender's mailbox for this service, or spaces if sender is provider
flag	text	1	application dependent
message	text	varies	

Fig. 5A

Description	Type	Length	Remarks
return mailbox	text	4	sender's mailbox for this service, or spaces if sender is provider
flag	text	1	"s" for subscription, or "c" for cancellation
task	binary	2	task number
user id	text	3	extracted by API
key	text	8	key

Fig. 5B

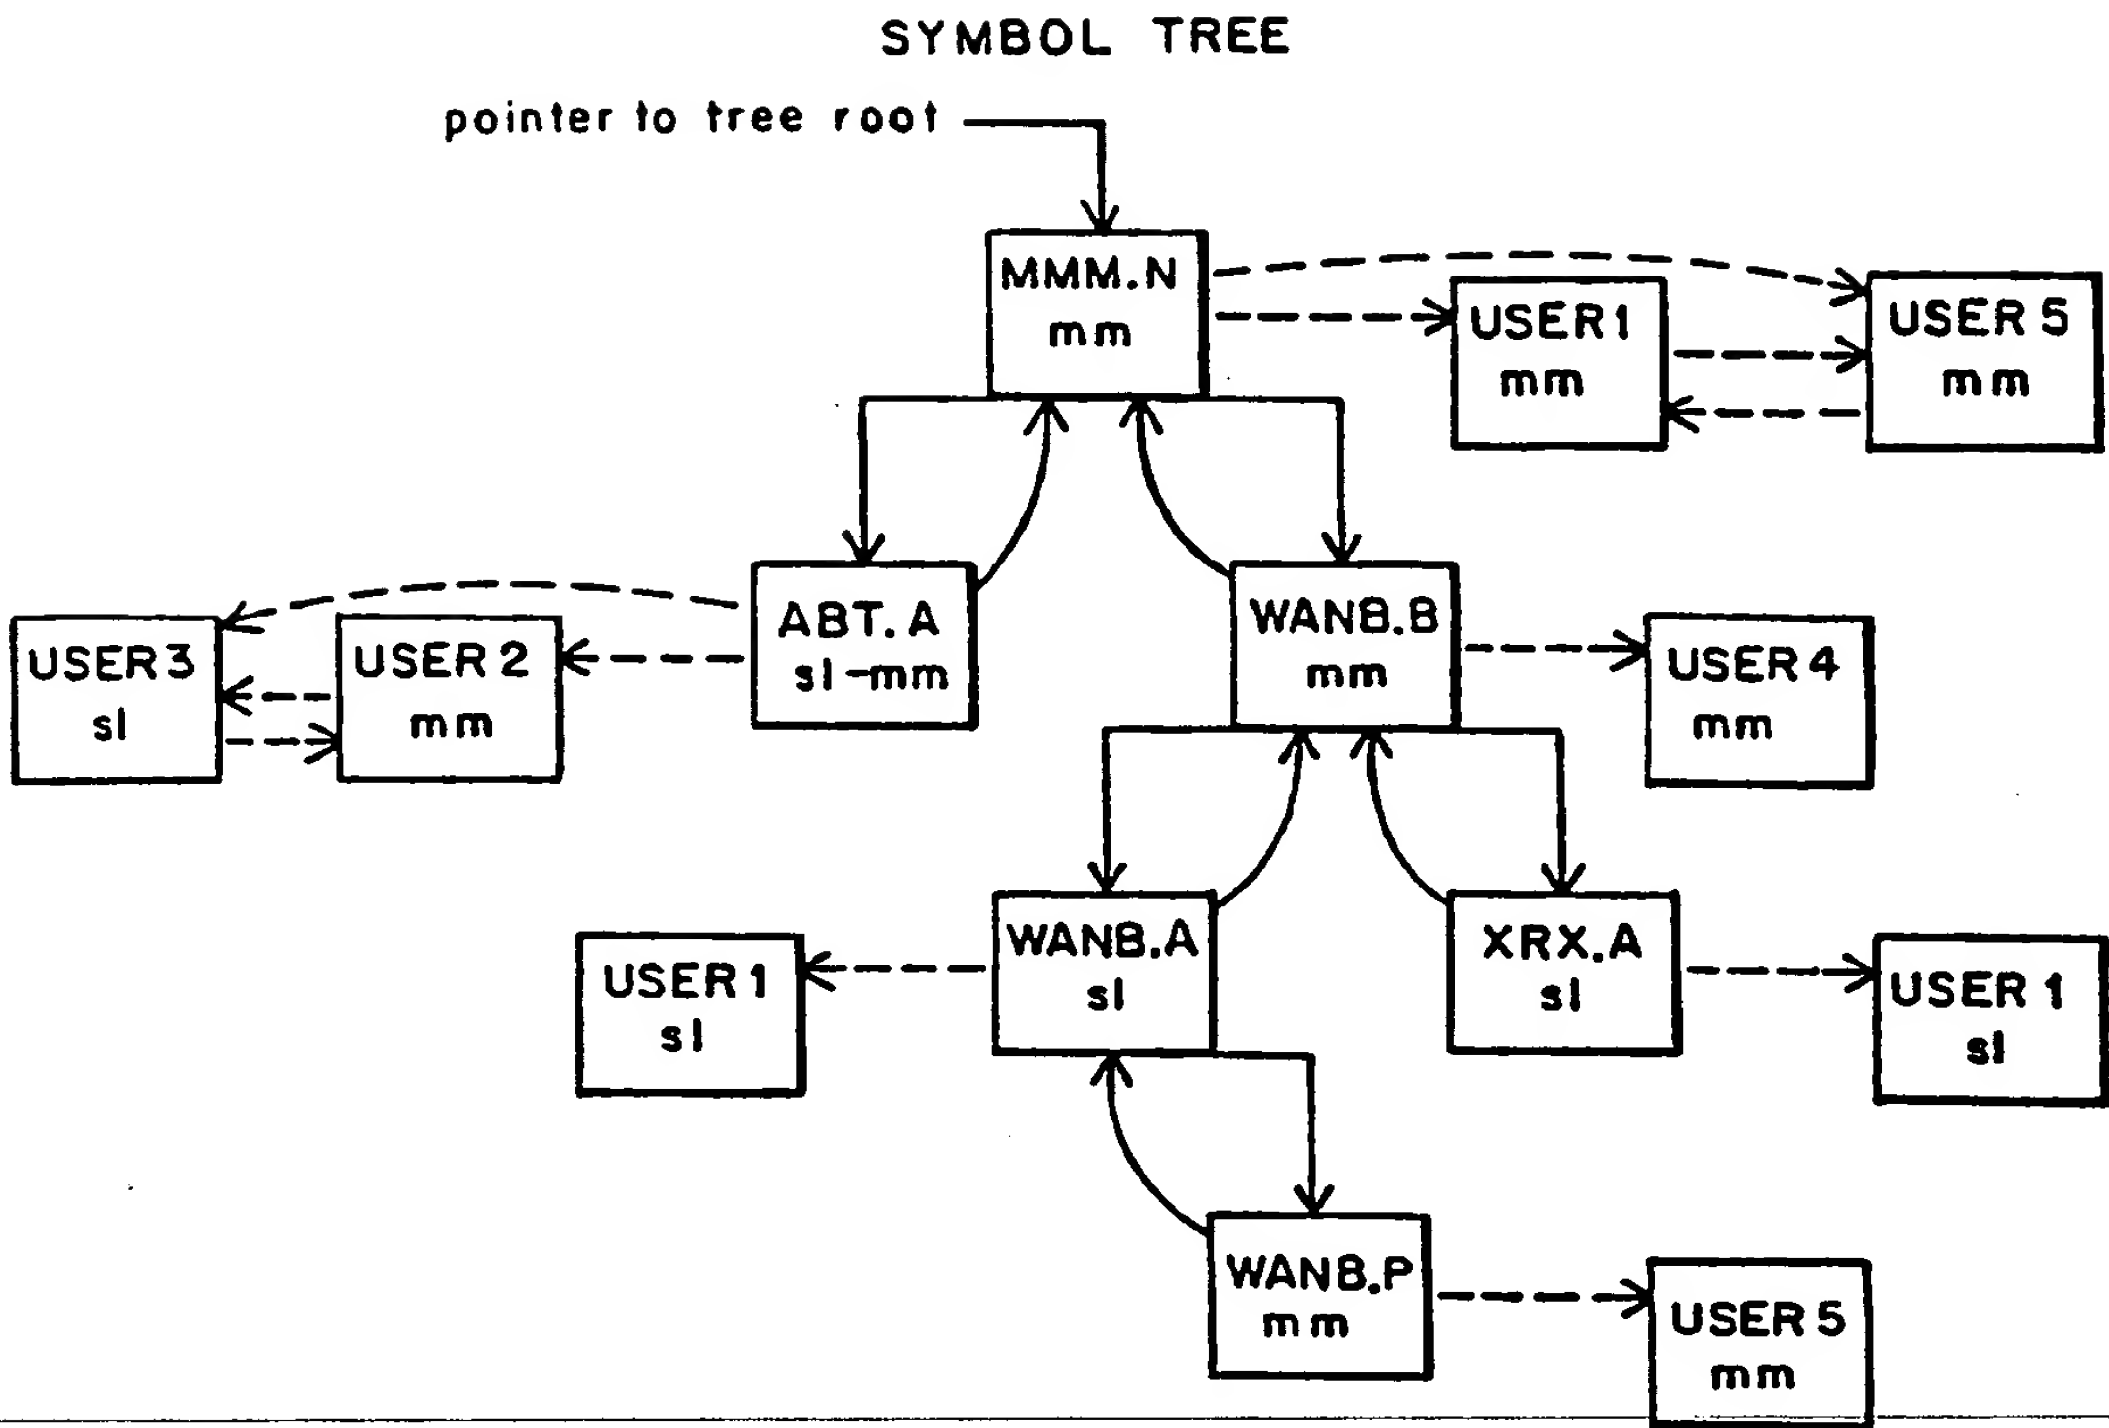
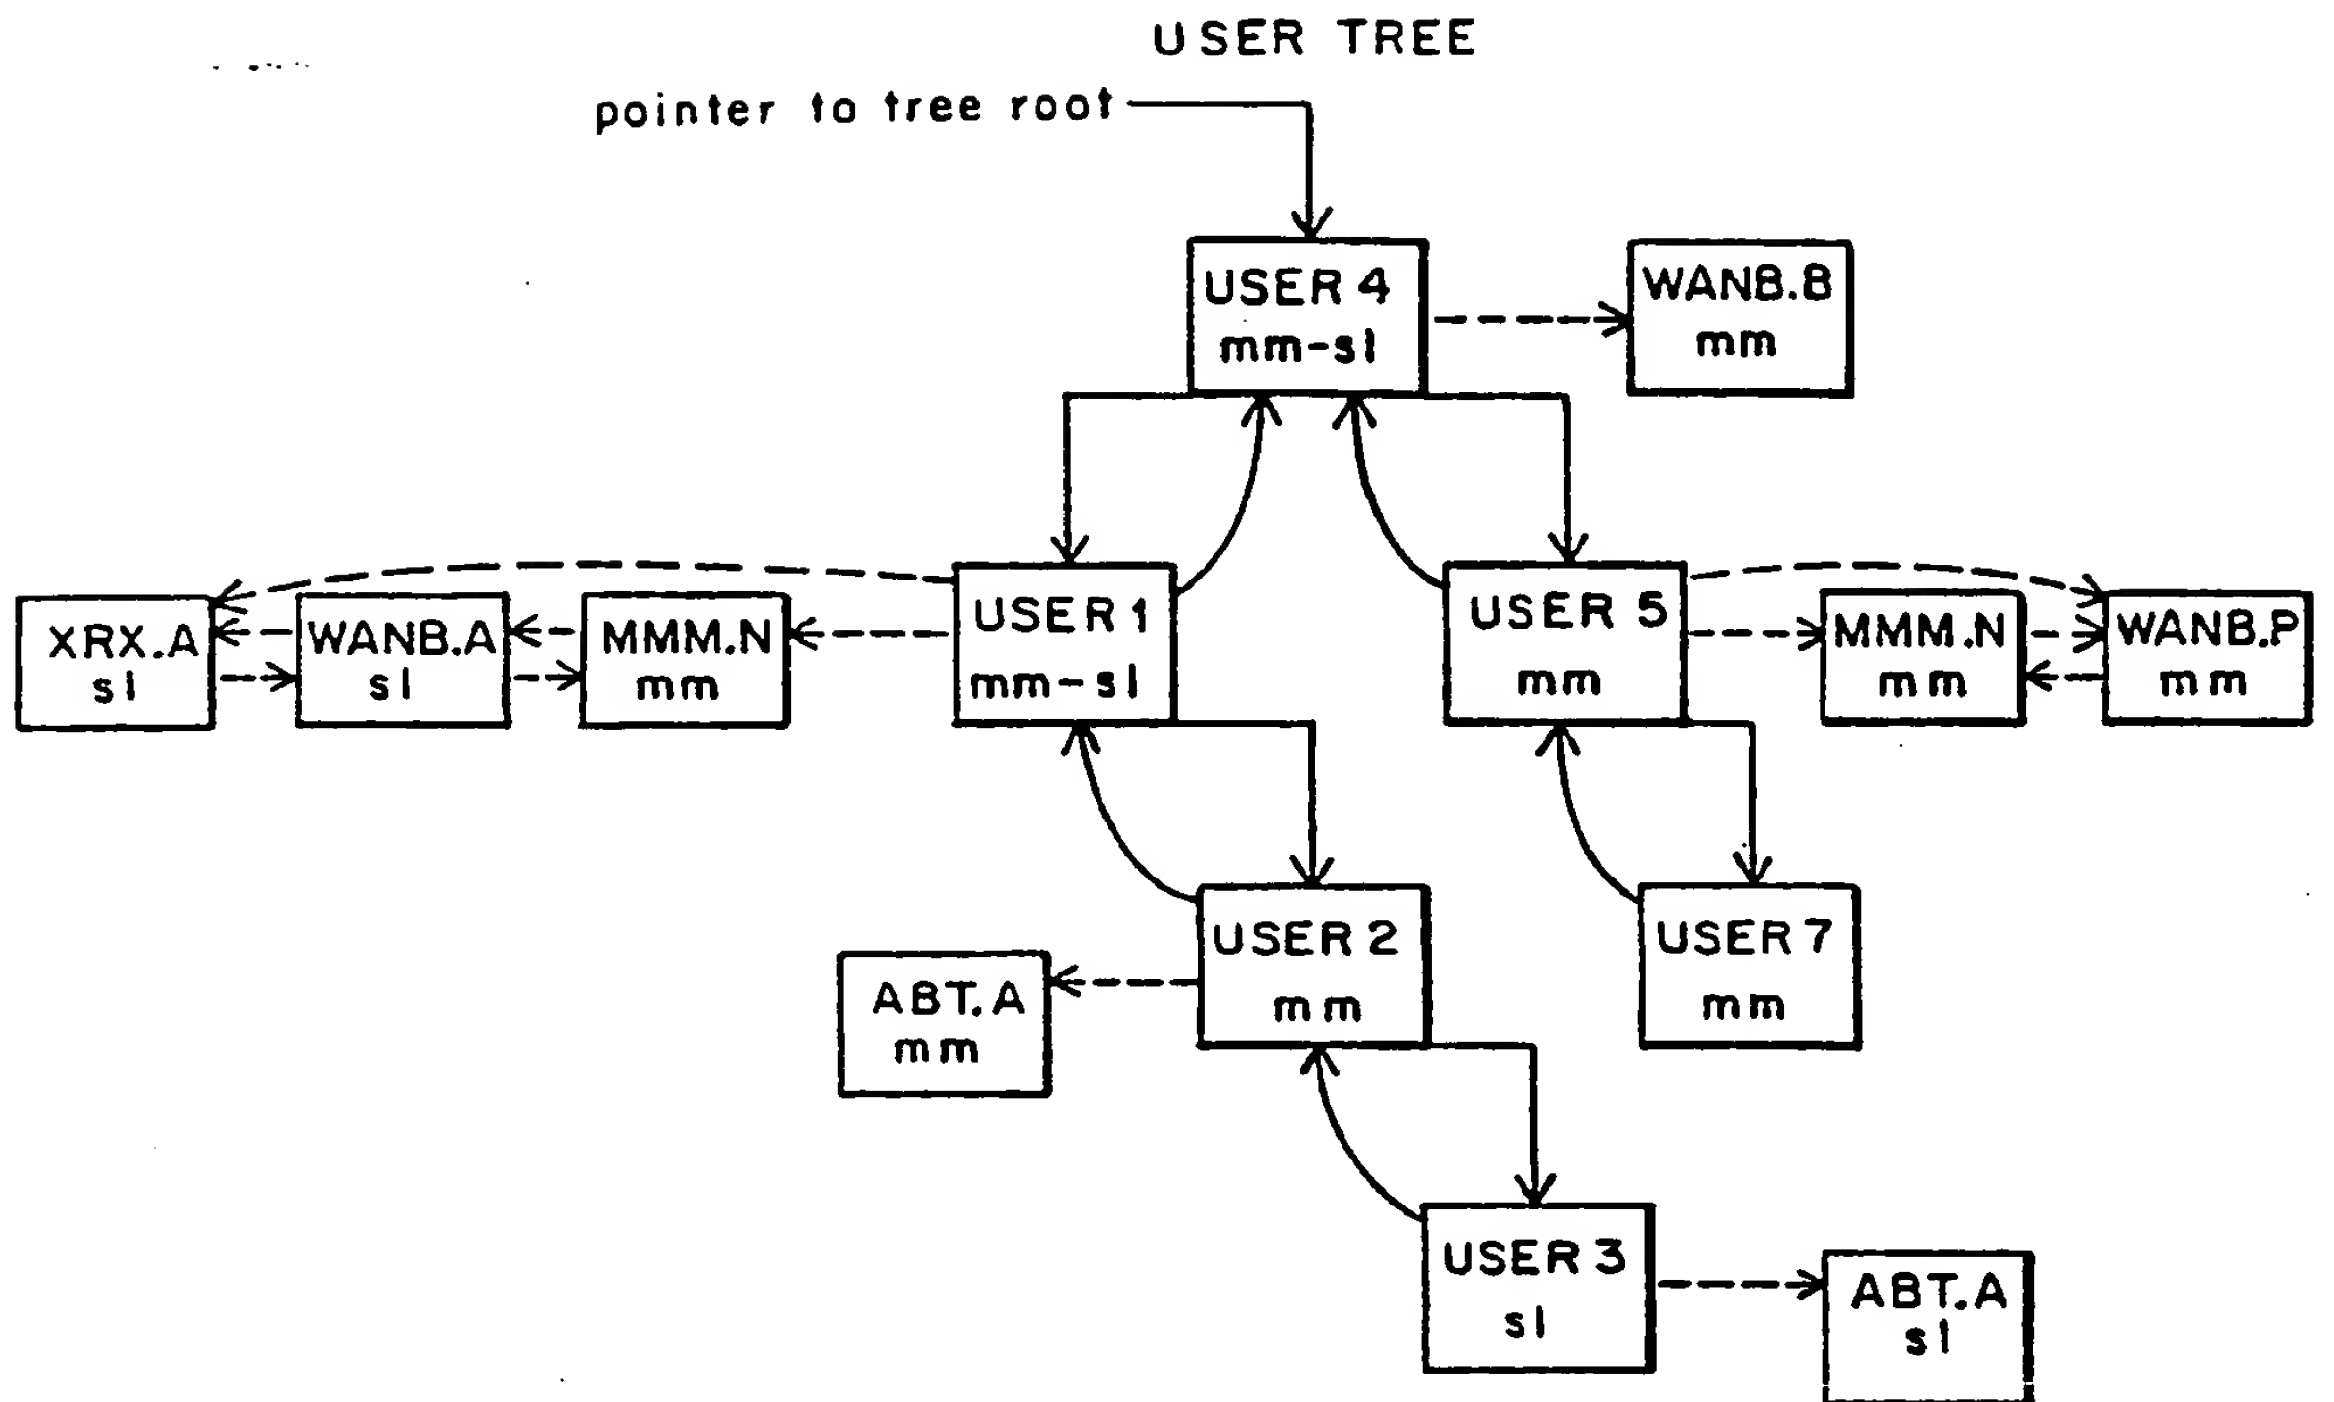
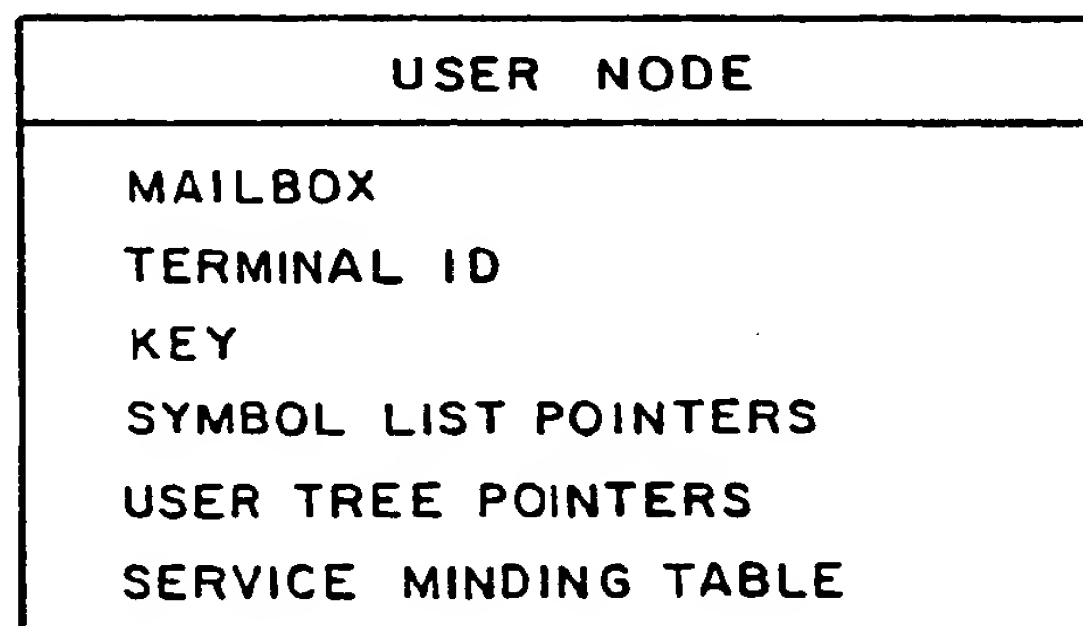
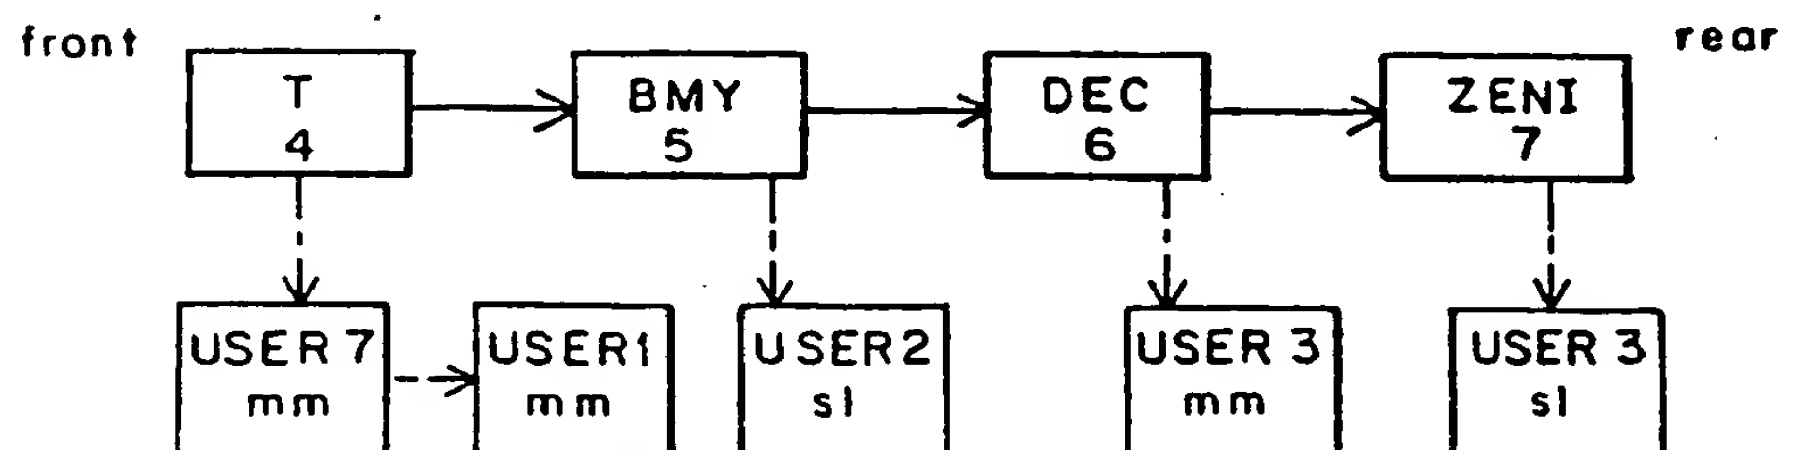


Fig. 6A

SYMBOL NODE
SECURITY SYMBOL
EXCHANGE
PRICE
USER LIST POINTERS
SYMBOL TREE POINTERS
SERVICE MINDING TABLE

Fig. 6B

*Fig. 7A**Fig. 7B**Fig. 8*

12

EUROPEAN PATENT APPLICATION

21 Application number: 87112735.3

51 Int. Cl.⁵: G06F 9/46

22 Date of filing: 01.09.87

30 Priority: 03.09.86 US 903495

Lowell, MA 01851(US)

43 Date of publication of application:
09.03.88 Bulletin 88/10

72 Inventor: Cross, Charles B.
 3 Bonnie Lane
 Billerica, MA 01821(US)
 Inventor: Moy, Diana Y.
 6 Old Farm Road
 Wayland, MA 01778(US)

84 Designated Contracting States:
BE DE FR GB

83 Date of deferred publication of the search report:
09.05.90 Bulletin 90/19

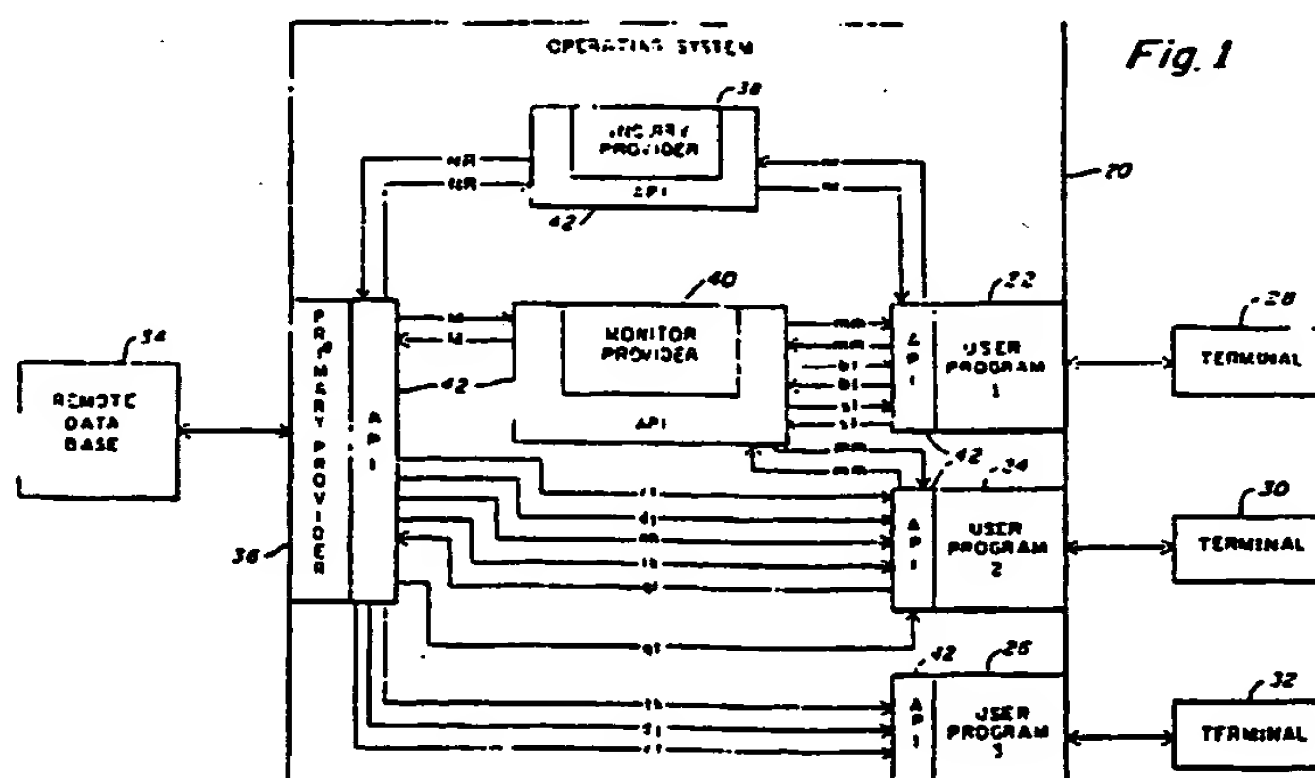
74 Representative: Behrens, Dieter, Dr.-Ing. et al
 Wuesthoff & Wuesthoff Patent- und
 Rechtsanwälte Schweigerstrasse 2
 D-8000 München 90(DE)

71 Applicant: WANG LABORATORIES INC.
One Industrial Avenue

54 Multitask subscription data retrieval system.

57 A multitask multiuser system provides for efficient transfer of data from a remote data base to individual subscribers and has particular utility in the distribution of data, especially of stock market data. A primary provider distributes the incoming data directly to user tasks or to an inquiry provider or a monitor provider. The inquiry provider responds to specific inquiries by users for information in the data base. The monitor provider maintains lists of information which are being monitored by the host computer for individual users. The inquiry provider

and the monitor provider do not repeat requests to the remote data base where a similar request is already pending from another user. Data transfer paths between tasks are established by a code module which may be linked to any of the tasks. The transfer paths are established using information from a configuration list and they are monitored by the operating system through a wait list established for each user task. Providers in the system may establish subscriber lists through the code module.





European Patent
Office

EUROPEAN SEARCH REPORT

Application Number

EP 87 11 2735

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int. Cl.4)
X	GB-A-2 161 003 (HIGGINS) * Figures 1-4; page 1, line 6 - page 7, line 125 *	4,13,14	G 06 F 9/46
A	---	1-3,5-12	
A	GB-A-1 489 573 (AUTOMATED REAL-TIME INVESTMENTS EXCHANGE LTD) * Figures 1,2; page 1, line 27 - page 2, line 2; page 39, line 45 - page 41, line 62 *	1-5	
A	---	1-5	
A	US-A-3 597 741 (SIERACKI) * Abstract *	1-5	
A	---	1-5	
A	LA RECHERCHE, no. 153, March 1984, pages 407-409; GARDARIN: "Bases de données; la percée des machines spécialisées" * Figures 1,2 *		

			TECHNICAL FIELDS SEARCHED (Int. Cl.4)
			G 06 F 9 G 06 F 15
The present search report has been drawn up for all claims			
Place of search THE HAGUE		Date of completion of the search 12-02-1990	Examiner SCHENKELS P.F.
CATEGORY OF CITED DOCUMENTS X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document			

EPO FORM 1503 03.82 (P0401)